J. Webster (ed.), *Wiley Encyclopedia of Electrical and Electronics Engineering*

HYPERMEDIA

In a general sense hypertext allows pieces of information to be connected to each other, and a reader is able to follow these connections directly. The connections are termed links and the action of following the links is termed navigation. The World Wide Web is the most widely accessible example of hypertext. Although it is a subset of the full hypertext model, it is an extremely powerful application. While *hypertext* is the original term, the term *hypermedia* is also in common use. Hypertext is not limited to use with text but can also include other media, such as images or even audio or video. Hypermedia often refers to this more general use of the term *hypertext* but can also indicate that there are links within a multimedia presentation.

When reading a hypertext document, some parts of the text are highlighted to indicate that the reader can select them. The most common method of interaction is to select them by clicking with a mouse pointing device. For example, in Fig. 1, the words representing the starting point of a link are underlined. When the reader wishes to follow a link, the mouse cursor is placed over the words (for example, "Research Interests"), and the reader clicks the mouse button. The display then changes to show the new information, shown in Fig. 2. The area where the reader clicks is an anchor and is commonly referred to as a hotspot. An accessible and complete overview of multimedia and hypotext is given in J. Nielsen (20).

History

The concept of hypertext has a very long history. Early religious works were a form of hypertext, in which scribes wrote comments on the original texts (these comments were referred to in later texts). These hypertexts were paper based. In 1945, before the advent of electronic computers, Vannevar Bush (1,2) proposed a mechanized system to implement the concept of hypertext.

> Consider a future device for individual use, which is a sort of mechanized private file and library. It needs a name, and to coin one at random, "memex" will do. A memex is a device in which an individual stores all his books, records, and communications, and which is mechanized so that it may be consulted with exceeding speed and flexibility. It is an enlarged intimate supplement to his memory. It affords an immediate step, however, to associative indexing, the basic idea of which is a provision whereby any item may be caused at will to select immediately and automatically another. This is the essential feature of the memex. The process of tying two items together is the important thing.

Douglas Englebart pioneered a system exhibiting a number of hypertext features, called *NLS* (online system) (3,4). This allowed researchers to create and share documents locally and remotely. NLS included an implementation of multiple document windows and the ability to refer directly to parts of the documents.

Theodor Nelson coined the term *hypertext,* and in particular he wanted to use the technology available at the time to support writers. He designed a system, called Xanadu, which would allow every letter of any

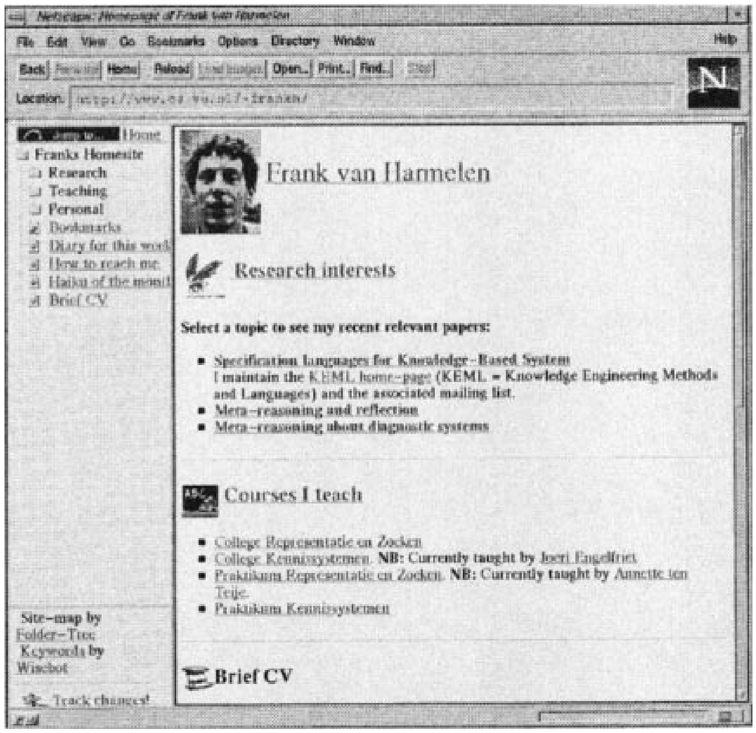

Fig. 1. Example hypertext screen 1.

document to be directly referenced or included in any other document while preserving its original author. Copyright protection was one of the underlying issues.

This gives a very brief view of some of the early hypertext visionaries. A large amount of research took place through the 1970s and 1980s, when different hypertext systems were implemented and used by limited numbers of people. This changed dramatically in the 1990s when the World Wide Web took off (5). Tim Berners-Lee demonstrated a text-only browser at the ACM Hypertext '91 conference. During the few years that followed, the use of the Internet changed dramatically, based on the two main components of the Web—the document language, Hypertext Markup Language (HTML) (6) and Hypertext Transfer Protocol, or HTTP protocol (7). HTML documents were the first globally accessible hypertext documents.

Applications

While hypertext is an enabling technology, and in essence domain neutral, there are properties of an application area that makes it particularly suited to hypertext. These are that the information can be partitioned into self-contained parts and that the topic is relatively complex, making cross references to several topics useful.

Medical knowledge is a characteristic example, since a part of the body is related in an anatomical way to neighboring parts but is also related in a functional way to other parts. For example, the lungs are close to the stomach but are part of a separate functional system. The underlying structure of the material can be reflected in the linking structure of the documents. Airplane and automobile manuals and collections of legal documents are similar types of examples.

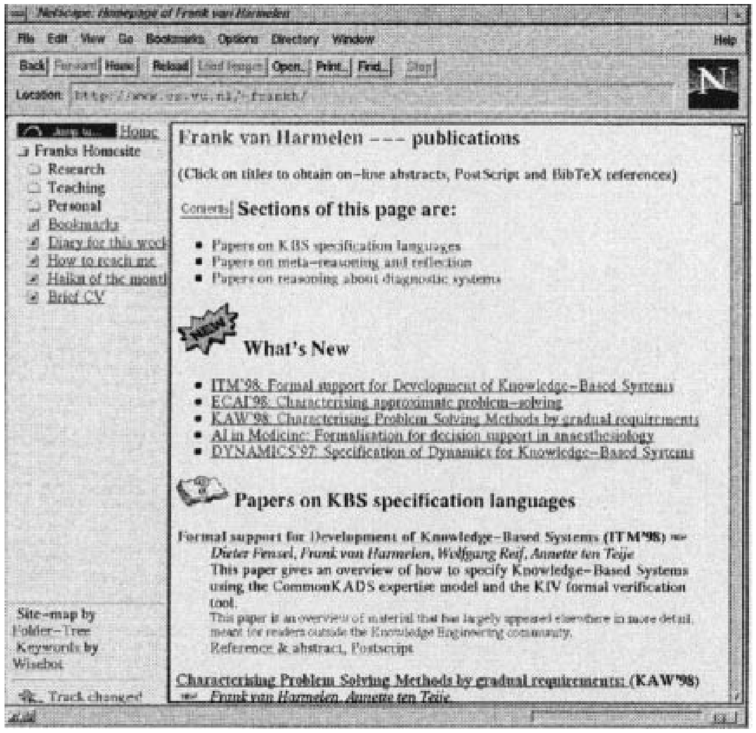

Fig. 2. Example hypertext screen 2.

On-line news is a slightly different type of example. Here there are topics that remain, more or less, constant, such as today's value of the New York Stock Exchange (*NYSE*). This can be linked to all the previous values but can also be linked to the values for today in the European and Asian markets. Hypertext allows the different connections to be explored by the reader, choosing the ones most relevant to the task.

We discuss different application areas of hypertext in more depth later in this article.

Definitions of Hypertext, Multimedia, and Hypermedia

Hypertext, multimedia, and hypermedia are commonly used terms with no consensus on their definitions, so we give the definitions as used throughout this article. A *hypertext document* is a collection of self-contained information units and referencing information, called links [Fig. 3(a)]. A *hypertext presentation* is the runtime manifestation of one or more hypertext documents with which a reader can interact. The information units in a hypertext document may include media types other than text. A commonly used term for this is *hypermedia.*

A *multimedia document* is a collection of information units and associated synchronization information [Fig. 3(b)]. A *multimedia presentation* is the runtime manifestation of a multimedia document. A reader can interact with a multimedia presentation by, for example, starting or pausing the presentation.

We use the term *hypermedia document* to denote a collection of information units along with referencing and synchronization information [Fig. 3(c)]. A hypermedia document is thus a collection of multimedia documents along with referencing information. A *hypermedia presentation* is the runtime manifestation of one or more hypermedia documents. A reader can interact with a hypermedia presentation either as a multimedia

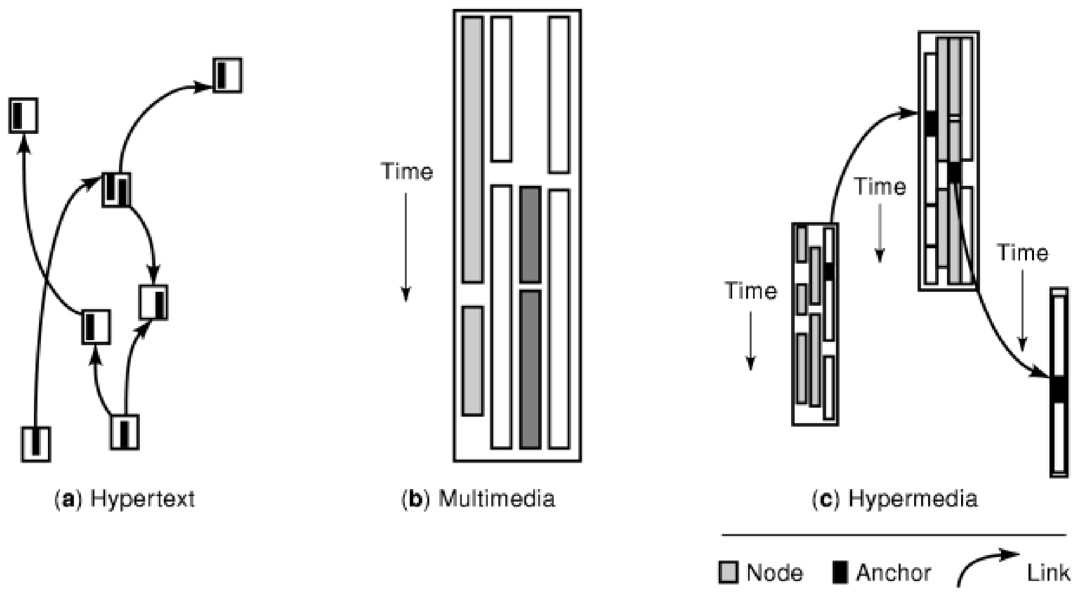

Fig. 3. Hypertext, multimedia, and hypermedia documents.

presentation, by starting or pausing a multimedia presentation, or as a hypertext presentation, by navigating through the information units.

Hypertext Systems

A hypertext system is the program for interpreting hypertext documents and presenting them to the reader. The system includes functionalities, such as displaying the text on the screen, highlighting the starting points for links, interpreting the reader's interactions with the document, and following the chosen links. A wide variety of hypertext systems have been built. To give a flavor of the diversity, we give some illustrations here.

KMS. Knowledge Management System (KMS) (1) is a frame-based hypertext system based on the research system ZOG (8). Frames can include text, graphics and links, where links are separated from the text and graphics and denoted by a marker. A three-button mouse is used extensively for interaction, and cursor feedback is used to explain the action of each of the mouse buttons. One of the design goals was to keep the response times of the system to under one second. This reduces user frustration and facilitates users in being able to reorient themselves by going back and forth in the document space. The structure of the linked frames is basically hierarchical, although crosslinks can be created.

Notecards. The Notecards system (9) manipulates notes that consist of, for example, text or graphics in separate windows. These notes can be filed in one or more file boxes. Parts of text, in the text nodes, can link to other windows. These are indicated by drawing a box around the appropriate text. A map of links among members of a set of notes can be generated. The system was designed to provide support for information analysts, with the intention of allowing analysts to express the conceptual models they form.

Intermedia. Intermedia's documents are text based in scrollable windows, which can include graphics (10). The links are separate from the text and graphics and are denoted by a marker in the window. A description of the link destination is attached to the marker by placing a description close to the marker or, in the case that the link has several destinations, a menu is displayed. Intermedia is one of the few hypertext systems that supports links with multiple destinations. Similar to NoteCards, Intermedia has a facility for displaying sets of links graphically. Because of the separation of links from endpoints, Intermedia can permit the same sets of information to use different or multiple sets of links. A set of links is called a web. Intermedia has been used to a great degree in teaching at Brown University.

Mosaic. The Mosaic browser (11) was developed at the US National Center for Supercomputing Applications (*NCSA*). Mosaic added an important factor to the World Wide Web—a graphical user interface to browsing HTML documents. End users could read text in a user-friendly environment and point at the links they wanted to follow. Pictures and icons could also be included within a text page. This was extended further, allowing parts of pictures to be associated with different link destinations. If other document types were not displayable by the browser directly, then an appropriate external viewer could be invoked. In addition to the display of HTML documents, the browser also provided navigation tools, allowing users to go back to documents they had seen before and to mark those of particular interest.

Another important success factor was that the browser was developed for the three common operating platforms: Microsoft Windows, Apple Macintosh, and X Windows. The Mosaic browser was developed commercially as Netscape.

Microcosm. The philosophy behind the Microcosm system (12) is to allow the creation of links among documents that are not part of a specialist hypertext authoring environment. This requires the specification of anchors and links externally to the documents being linked. The Microcosm designers provide links without forcing the author to create each one separately by hand. They also provide the facility for creating links from any occurrence of a word without requiring the author to specify its position in the document (in other words, without forcing the author to specify the anchors individually). Microcosm also allows for the creation of linkbases. A linkbase is collection of links. Several linkbases can apply to the same sets of documents (for example, for different authors or for different groups of readers).

Application Areas

While hypertext can be used for a broad range of purposes, a number of areas have made particular use of hypertext. These include computer-based learning, design histories, collaboration in writing, technical documentation, and entertainment.

Computer-Based Learning. Study material can be provided to students as hypertext by just copying a text book. To make use of the advantages of hypertext, however, there has to be something more (for example, by allowing students each to have their own view of the information and allowing them to add their own comments). These can then be shared with a tutor or with their peers. Literature studies are a typical example of this type of use, where there is never a “correct” answer, and students can create their own interpretations and compare them with those of their peers.

Another advantage of learning from the computer is that simulations of the ideas being explained (for example, the laws of gravity or the effects of heat loss from a house), can be illustrated directly by linking to a simulation. In medical domains, for example, X rays and videos of joint motion can be included in the material. The hypertext is already on line, so that the environment can be extended to record which material students have already seen, keep scores on tests, or even to prevent students from seeing material until they have read sufficient preparatory material.

Collaborative Writing. Writing is not always a solitary process, and often collaboration is carried out with pen and paper. This is not always satisfactory, since paper cannot be stretched to accommodate long comments. Hypertext offers facilities to overcome this problem by allowing links to information inserted by a colleague.

A number of the early hypertext systems were designed explicitly for collaborative writing, NLS in particular. This allowed researchers to speak to each other, see each other via video, point to the same documents on their screens, and type additional material into their documents. Other systems allowed writers to produce text, pass it on to colleagues, have each colleague comment on it, and then pass all the comments back again.

Design Histories. Hypertext systems are created by software engineers, and software engineers soon noticed that they could use hypertext for their own purposes. In particular, a number of design decisions are made in the course of creating a software system, and these decisions can be recorded in a hypertext system. For example, a proposition is stated, and linked to it are all the arguments for, all the arguments against, and the decision that was taken. Later in the software life cycle the decision can be reanalyzed and if, for example, an assumption turned out to be incorrect, the decision can be reversed—only after recording the new line of reasoning, of course. The gIBIS system (1) implemented these ideas and uses a graphical interface for illustrating the rhetoric of the argumentation.

Technical Documentation. Complex technical documentation printed on paper is both very heavy and provides no easy means of finding what the reader wants. By putting the information on line and turning cross references into hypertext links, the reader's life becomes much easier. To start with, readers can use an index or a search mechanism for finding the topic they are looking for. When they have found the right "page," they can browse the links for explanations of, for example, terms they are not familiar with.

Car and airplane maintenance are typical examples of this type of use. The added advantage of having the documentation on line is that the computer can be used for even richer illustrations. For example, if the reader wishes to change a spark plug, then the documentation can include a video with spoken commentary showing exactly where the components can be found and how they can be disassembled. Simulations can also be run on the computer to illustrate certain points in detail. As well as providing information, the computer can control information. For example, in a chemical plant or a power station the documentation can be directly connected with the software that runs the plant. This allows experienced controllers to use the documentation while running the plant and allows trainees to learn from the documentation while seeing the live values of parameters from the plant.

An important aspect of technical documentation is that the artifact that it is describing is likely to outlive a number of generations of computers. This means that the information has to be stored in a system-independent way. Documents are often described using Standard Generalized Mark-up Language (SGML) (13).

Entertainment. Hypertext offers the potential of a new art form. While paper "hypertexts" have been created, they tend to be cumbersome to read, since the reader has to "skip to page X" at every choice point in the story. By putting these online, a writer can offer the reader not only an imaginary world, but multiple worlds in which the reader can influence the progress of the "story." There is no longer a single story, but threads of intertwined stories.

In the film world, experiments have been done in which the same action has been shot from the perspective of different characters. The multiple streams are shown simultaneously (say on different television channels). A viewer can choose which character to follow. The difference with hypertext is that the viewer cannot see everything but must make a choice.

Games can also be created using hyperlinks. For example, a world can be created and the reader can explore the world by following hyperlinks. In no time the reader can become completely disoriented, but in contrast to the technical applications, this was precisely the intention of the author.

Again, the applications can go beyond simple hypertext and include videos and animation. Adventure games can be made in which players can interact with the environment (for example, draw in the environment, or pick up objects to take with them).

Visual Design for Anchors

Paper-based documents have developed their own visual conventions during their extended existence. When creating hypermedia documents, extra information needs to be expressed to readers so they can identify

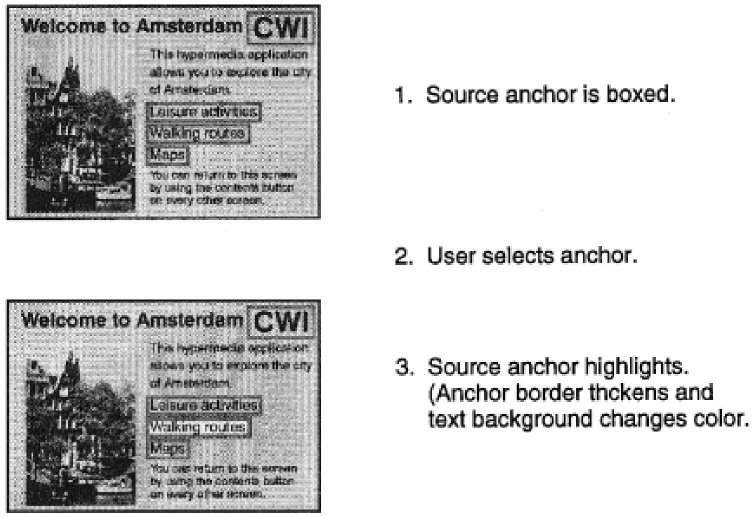

Fig. 4. Anchor and transition styles on following a link.

potential links and interact with them successfully. A powerful method of expressing this is visually. The problem faced by hypermedia designers is that they need to introduce new conventions, which may violate existing conventions or just add ugly clutter to the screen.

Style information for an anchor is needed for specifying how the visual, or even audible, characteristics of an anchor can be specified so that a user is aware that an anchor is indicating the start of a link. Examples of anchor styles are as follows:

- There is a border around the anchor value, as illustrated in Fig. 4.
- There is a small icon next to the anchor value.
- For text items, different color or styles, such as underline or italic, is used.
- The anchor value changes appearance (e.g., color) when the mouse cursor is over it.
- The mouse cursor shape changes when it is over the anchor value.

When a user selects an anchor to follow a link, there may also be style information associated with this action. For example, the source anchor may highlight (to acknowledge that the action has been registered) before the destination of the link is displayed. The destination anchor may also be highlighted briefly to distinguish it from any other anchors present in the destination. The following are examples of anchor highlight styles:

- The appearance (e.g., width or color of the border) changes, as illustrated in Fig. 4.
- The appearance of the icon changes.
- For text items, the style and/or color changes.
- The anchor value flashes.
- The anchor value changes color.

The style of a source anchor may depend on other properties of the link emanating from it (for example, whether the reader has already seen the destination of the link). Further information on screen design for hypertext can be found in Ref. 19.

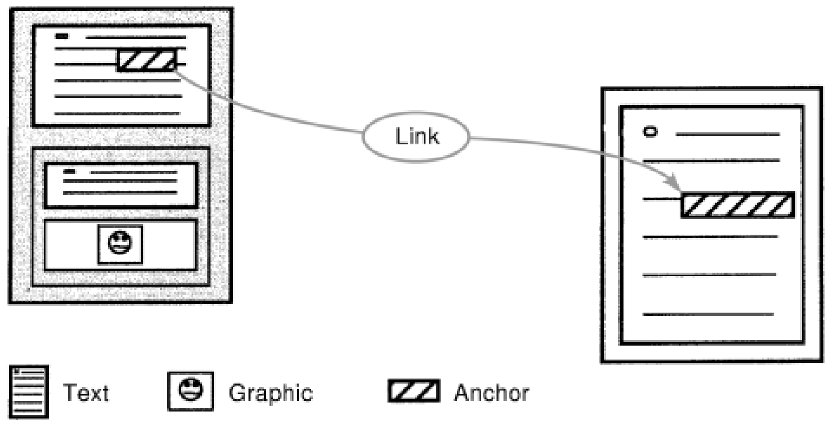

Fig. 5. Dexter hypertext model overview. A composite component (left) is linked to an atomic component (right).

Document Models

While the user view is a useful interaction paradigm, in order to make full use of hypertext in a powerful publishing environment, we need to model hypermedia documents for processing. This allows information to be reused in multiple systems and also allows the same information to be reused multiple times within a single system. This work borders on that of electronic publishing using technologies such as SGML (13) and HyTime (14).

Dexter. The Dexter hypertext reference model (15) was developed as a reference model to rationalize and make explicit the concepts embedded in the then existing hypertext systems. The Dexter model divides a hypertext system into three layers: a *within-component layer,* where the details of the content and internal structure of the different media items are stored; the *storage layer,* where the hypertext structure is stored; and the *runtime layer,* where information used for presenting the hypertext is stored and user interaction is handled. The Dexter model describes the storage layer in detail, and it is this layer that is most relevant to a hypermedia model and of which we give a brief description here.

The Dexter model introduces atomic, composite, and link components and anchors. Atomic and composite components are related to each other via link components, where anchors specify the location of the ends of the links. This is shown schematically in Fig. 5.

Each component has its own unique identifier. A reference to a component can be made directly to its unique identifier or via a more general component specification. The latter requires a resolver function to “resolve” it to a unique identifier (for instance, to allow the addressing of a component by means of an SQL [Structured Query Language] database query).

Atomic Component. An *atomic component* contains four parts—presentation specification, attributes, a list of anchors, and content.

- The *presentation specification* holds a description of how the component should be displayed by the system.
- The *attributes* allow a semantic description of the component to be recorded.
- An *anchor* is composed of an anchor identifier and a data-dependent anchor value. The anchor identifier is unique within a component and allows the anchor to be referred to from a link component. The anchor value specifies a part of the content of the atomic component and is the only place in the model where the data type of the content is required. The anchor value is used as the base of the hotspot in the presentation of the document.
- The *content* is a media item of a single data type.

Composite Component. A *composite component* is a collection of other components (atomic, composite, or link) that can be treated as a single component. Its structure is the same as an atomic component with, in addition, a list of child components. The structuring of the components is restricted to a directed, acyclic graph. The anchors of a composite component refer to the content of that component.

Link. A *link* is a connection among two or more components. Its structure is the same as an atomic component, with a list of specifiers replacing the content. A *specifier* defines an endpoint of the link. It consists of an anchor, a direction, and a presentation specification.

A single link component can allow the expression of a range of link complexities, including a simple one-source, one-destination, uni-directional link (for example, links in HTML) or a far more complex multisource, multidestination, bidirectional link.

Amsterdam Hypermedia Model. While the Dexter model is adequate for describing hypertext, it does not incorporate the temporal aspects of hypermedia. To provide such a model for hypermedia, four fundamental types of relationships must be described in addition to the media items included in the presentation: structural, timing, layout, and interaction relations. Structural relationships define logical connections among items, including the grouping of items to be displayed together and the specification of links among these groupings. Timing relations specify the presentation dependencies among media items, possibly stored at different sites. Layout specifications state where screen-based media are to be sized and placed, either in relation to each other or to the presentation as a whole. Interactions, through navigation, give the end user the choice of jumping to related information.

The Amsterdam Hypermedia Model (*AHM*) (16) describes a model for structured, multimedia information. A diagrammatic impression of the AHM is given in Fig. 6.

Timing Information. Timing information is needed to express when each of the media elements composing the presentation will appear on (and disappear from) the screen. This information can be given by specifying a time when an element should appear relative to the complete presentation (e.g., 30 s after the beginning of the presentation, a subtitle corresponding to the spoken commentary should appear) or relative to other elements being played in the presentation (e.g., 3.5 s after the spoken commentary begins, the subtitle should appear). The AHM allows the specification of timing relations defined between single items, groups of media items, or between a single item and a group. These timing relations are specified in the model as *synchronization arcs.* These can be used to give exact timing relations but can also be used to specify more flexible constraints, such as “play the second media item 3 ± 0.3 s after the first.”

Structural Information. The structure items of the model are described here briefly.

- *Composition* plays an important role in multimedia presentations, where almost all presentations contain more than one media element. A composition structure allows a group of media items to be created that can then be treated as a single object. For example, a company logo can be grouped with a spoken commentary and included in several places in the overall presentation. The composite structure is used to store the timing information specified among its constituent elements.
- *Anchors* are a means of indexing into the data of a media item, allowing a part of the item to be referred to in a media-independent manner. The data specified by an anchor can be used at the beginning or end of a link. For example, within text, an anchor might define a character string; in an image, an area on the screen. In continuous media, such as video or animation, the area on the screen may change with time. Indeed, the hotspot may appear only for part of the duration of the media item—for example, the video item in Fig. 6 shows two hotspots, which are displayed at different times. One of the uses of a dynamic hotspot is to follow moving objects within a video or animation. In a composite item (for example, the scene represented by Fig. 6), an anchor can refer to the scene as a whole (allowing navigation to the complete scene) or indirectly to anchors in the media items. For example, a video of a bouncing ball may be accompanied by a text about the ball. Both the moving image of the ball and the word *ball* in the text are part of the same (composite)

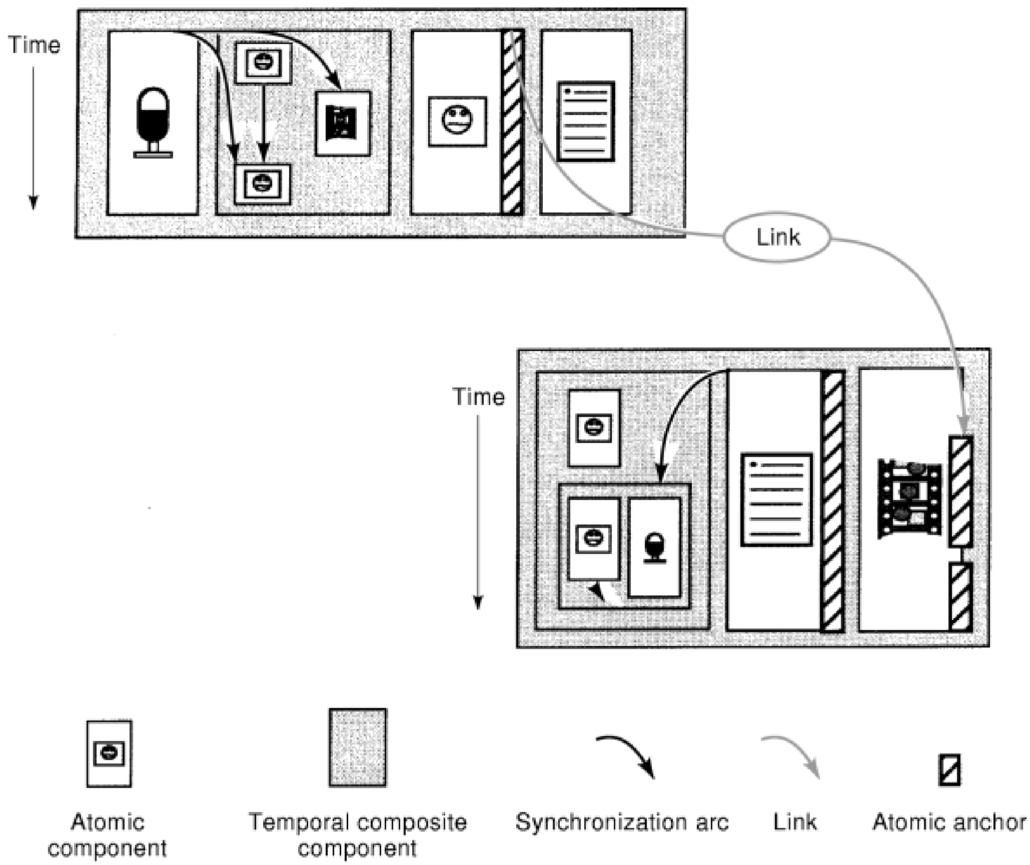

Fig. 6. Amsterdam hypermedia model overview.

anchor. The anchor in the composite representing the concept "ball" refers to two anchors—the anchor in the text item and one of the anchors in the video item. When the author wishes to relate other information to the ball's description (e.g., how gravity affects bouncing), then a link can be made to the composite "ball" anchor.

- *Links* enable the end user to jump to information deemed by the author to be relevant to the current presentation. The link can lead to another presentation or a different part of the same presentation. *Link context* is needed for specifying which media items on the screen are affected when a link is followed. For example, in a presentation composed of multiple media items, the complete presentation may be replaced or only a part. There is also a choice of whether the information at the destination of the link replaces the scene that was being played, or whether it is displayed in addition. When the current presentation also contains one or more continuous media items, there is a further choice as to whether the current scene should stop or continue playing when the link is followed.

Layout Information. Layout information includes the size of items displayed in a presentation and their position on the screen. Information about styles, such as font typefaces and styles, is also required. This information can be attached to individual media items, but in the model this is done through the use of *channels.* This is particularly useful, for example, when a similar layout is required but the sizes of media items vary. Each media item is then scaled to fit into the assigned channel.

Expressing the Model. The AHM does not prescribe a language for specifying a presentation conforming to the model; it could be expressed in a system-independent language such as the HyTime (Hypermedia/Time-Based Document Structuring Language) international standard (14) based on SGML (13). Similarly, transmission of (groups of) multimedia objects conforming to the model could be carried out using a standard such as *MHEG* (Multimedia and Hypermedia Expert Group) (17). SMIL (Synchronized Multimedia Integration Language) (18), which was developed for the World Wide Web, is based on the same concepts underlying the AHM.

Future Directions

Looking at the World Wide Web (the Web), hypertext appears to be already "solved": users can publish their own documents and access documents from Internet sites around the globe, and links can be made within and between any of these documents. There is, however, still plenty of work to be done and we discuss here a number of directions that we see as relevant in the near future of hypermedia documents: extensible document languages, multimedia, open hypermedia systems and metadata.

Extensible Document Languages. HTML is a specific language supported as part of the Web, so that users are constrained by the features it offers. Work has been carried out to allow different document languages to be specified using the Extensible Markup Language (*XML*) (22). This allows content creators to define their own document types and yet have access to standard tools and playback environments. While this already allows an extra level of flexibility, the XML namespaces initiative (22) will go further and allow a single document to mix and match terms from different document specifications.

Multimedia. Multimedia has existed in the CD-ROM world for nearly a decade. The introduction of SMIL (18) has made multimedia documents on the Web possible. Since SMIL is an XML application, this means that SMIL expressions can be used, via XML namespaces, in other XML documents. Any XML document would be able, for instance, to incorporate links, where the destination of the link varies depending on when the reader interacts with it. The difference between using SMIL as a vehicle for multimedia rather than a more programming-based construction is that the document can be manipulated by all the standard XML tools, allowing, for example, styles (as specified using Extensible Style Language [22]) to be applied in the same way that they can be applied to hypertext documents.

Multimedia, and in particular temporal layout, will become an integral part of all documents, rather than something special that needs to be treated separately. Just as, for example, many word processing packages allow the creation of images.

Open Hypermedia Systems. Although the Web appears to be an open hypermedia system (*OHS*), it suffers because all the information available on the Web has to conform to the standards and conventions of the Web. An author cannot, for example, make a word processor document available in its native format, have parts of it link to other documents, and expect other Web users to be able to view it. The open hypermedia initiative (21) is directed at enabling hypermedia functionality for existing applications and components without forcing the material to a particular document format.

Using the services of an OHS, existing applications in the computing environment can become "hypermedia enabled," thus supporting linking to and from information managed by the application without altering the information itself. Problems that need to be overcome are that an application that is OHS enabled should be enabled for all OHSs; hypermedia structures created by one OHS should be understood by all other OHSs; and single links should, for example, be able to have one end in one environment and the other in another. For example, a link from a cell in a spreadsheet should be able to link to a date in a diary application.

In other words, what we are now used to on the Web should become part of the desktop computing environment, where links and anchors are not just part of a particular hypertext browser or editor, but are ubiquitous.

Metadata. Media objects, such as video or audio, can be made more accessible by associating with them descriptions of their content. Additional information about a document, such as author or creation date, is also useful. Both these types of information are metadata. Metadata can be useful for searching on, for example when looking for pictures on a specific topic, or for processing in a more general way, for example for validating a document.

The document formats HTML and SMIL include constructs for providing metadata about the document. Because they are part of the Web "product mix," they can also include Resource Description Framework (*RDF*) (22) information, which can be used to label the roles of document elements directly.

Integration. The extensions noted previously will not occur in isolation from one another, but will develop in parallel and in combinations. Integration of all extensions in a single (complex) environment will soon be possible.

An example application that involves mixing document languages, multimedia, and metadata is the flexible, or adaptive, hypermedia. These are presentations that alter depending on the user's preferences, prior knowledge, the platform's capabilities, or the available network bandwidth. From a single source document, decisions can be made as to which parts of the document should be displayed to the reader, or even generated. Different presentation styles are suited for different end-user platforms, mixing of document types is useful for richness of choice of type of presentation at runtime, multimedia allows temporal ordering of document elements and metadata can be used for choosing among multiple document elements.

However far these developments will bring us in the short term, it is likely to be another 50 years before the visions of Bush, Englebart, and Nelson (2,3,4) will be realized in full.

BIBLIOGRAPHY

1. J. Conklin Hypertext: An introduction and survey, *IEEE Comput.*, **20** (9): 17–41, 1987.
2. V. Bush As We May Think, *The Atlantic Monthly*, **176**: 101–108, 1945. [Online], http://www.w3.org/History/1945/vbush/ http://ww.isg.sfu.ca/~duchier/misc/vbush/
3. D. C. Engelbart A conceptual framework for the augmentation of man's intellect, in Howerton and Weeks (eds.), *Vistas in Information Handling*, *Organization and Groupware*, Washington, DC: Spartan Books, 1963, pp. 1–29. Republished in Irene Greif (ed.), Computer Supported Cooperative Work: A Book of Readings, San Mateo, CA: Morgan Kaufmann, 1988, pp. 35–65. Also republished in T. Nishigaki (ed.), NTT Publishing, 1992.
4. D. C. Engelbart W. K. English A research center for augmenting human intellect, *AFIPS Conf. Proc.*, **33** (1): pp. 395–410. [Online], 1968. Available http://www2.bootstrap.org/
5. R. Cailliau A Little History of the World Wide Web [Online], 1995. Available http://www.w3.org/History.html
6. D. Raggett A. Le Hors I. Jacobs HTML 4.0 Specification [Online], 1998. Available http://www.w3.org/TR/REC-html40/
7. H. Frystyk Nielsen J. Gettys HTTP—Hypertext Transfer Protocol [Online], 1998. Available: http://www.w3.org/Protocols/
8. D. McCracken R. M. Akscyn Experience with the ZOG human-computer interface system, *Int. J. Man-Mach. Stud.*, **21**: 293–310, 1984.
9. F. G. Halasz T. P. Moran R. H. Trigg NoteCards in a Nutshell, *ACM Conf. Human Factors Comput. Syst.*, Toronto, Canada, 45–52, 1987.
10. B. J. Haan *et al.* IRIS hypermedia services, *Commun. ACM*, **35** (1): 36–51, 1992.
11. B. R. Schatz J. B. Hardin NCSA Mosaic and the World Wide Web: Global hypermedia protocols for the Internet, *Science*, **265**: 895–901, 1994.
12. W. Hall H. Davis G. Hutchings *Rethinking Hypermedia: The Microcosm Approach*, Dordrecht, The Netherlands: Kluwer, 1996.

13. ISO, SGML. *Standard Generalized Markup Language*, ISO/IEC IS 8879: 1985, 1985.
14. ISO, *HyTime. Hypermedia / Time-based structuring language*, ISO/IEC 10744, 1997.
15. F. Halasz M. Schwartz The Dexter Hypertext Reference Model, *Commun. ACM*, **37** (2): 30–39, 1994.
16. L. Hardman D. C. A. Bulterman G. van Rossum The Amsterdam Hypermedia Model: Adding time and context to the Dexter Model, *Commun. ACM*, **37** (2): 50–62, 1994.
17. *ISO, MHEG* Part 5, ISO/IEC IS 13522-5, 1997.
18. P. Hoschka (ed.) Synchronized Multimedia Integration Language, W3C Proposed Recommendation. Authors: S. Bugaj et al., [Online], 1998. Available http://www.w3.org/TR/REC-smil.
19. P. Kahn K. Lenk Screen typography: Applying lessons of print to computer displays, *Seybold Report on Desktop Publishing*, **7** (11): 3–15, 1993.
20. J. Nielsen *Multimedia and Hypertext, The Internet and Beyond*, Boston: AP Professional, 1995.
21. U. K. Wiil Open hypermedia systems working group message from the OHSWG chair, [Online], 1997. Available http://www.csdl.tamu.edu/ohs/intro/chair.html.
22. World Wide Web consortium, Technical reports and publications, [Online], 1998. Available http://www.w3.org/TR/.

LYNDA HARDMAN
DICK C. A. BULTERMAN
CWI